

REDD Roads Rising?

examining spatial variation in causal impacts of infrastructure on deforestation

Alexander Pfaff[⊙] (alex.pfaff@duke.edu), presenting author

&

Juan Robalino[⊙] (robalino@catie.ac.cr)

Abstract *

We examine whether the deforestation impact of a new road varies by prior development. A focus on differences from the average impact is motivated by the potential for REDD, i.e. reduced emissions from deforestation and degradation that can earn global payments if a country chooses to spatially adjust its development investments in transport networks. Also, the way impacts vary helps to explain the previous lack of clarity on roads' effects. Causal evidence results from not only using lagged roads investments but also applying matching methods to Brazilian Amazon deforestation during periods from 1976 to 2004. We find that prior development (here proxied by distance to prior roads) affects baseline deforestation rates without new roads as well as the deforestation impact of a new road. Analysis not correcting for this easily generates spurious results, including sign reversals. Specifically, first-decade impact of a new road is relatively low if prior roads are close or far (pristine areas). In between those bounds, new roads significantly raise clearing. The non-monotonicity of development's effect on first-decade impact highlights dynamics.

Keywords

deforestation, roads, infrastructure, matching, Brazil, Amazon

[⊙] we share lead authorship

* This draft was made possible by a large group that includes Diego Herrera, Robert Walker and Eustaquio Reis. For financial support we thank two NASA LBA projects (on both of which Walker and Reis were the US and Brazilian PIs, respectively), the InterAmerican Institute (IAI) and the Earth Institute which employed both Pfaff and Robalino. For helpful comments, we thank audiences at IFPRI and at a number of LBA conferences over a number of years.

1. Introduction

Investment in new roads that cross forested frontiers have generated controversy over decades in part due to the apparent tradeoffs between development and environment and natural resources. Investing countries typically wish to expand employment and income. However, the migration and the production that result are often seen to have both local and global negative externalities.

At the moment, the greenhouse-gas emissions associated with deforestation in the tropics (in the temperate zones many forested areas are net sinks) are receiving great attention for a number of reasons. First, it is thought that these emissions could be reduced at relatively low cost. Second, should global payments for reduced emissions from deforestation and degradation (or ‘REDD’) be a mechanism for reducing those emissions, it could in principle bring net benefits to the host countries and more countries appear to be considering positively the potential for such transfers. Third, while any such contracts will be based on estimated impacts on climate change, frontier forest conservation is also thought to bring benefits for species and for particular stakeholders.

How much new roads drive tropical deforestation has long been discussed. On the one hand, the general assumption still seems to be that such investments in transport will raise deforestation. On the other hand, even after a second wave of empirical work within countries (the first wave involved cross-country regressions (more below)), the impacts of roads relative to other factors have not always clearly corresponded to the magnitudes that would appear to be assumed. There are even settings involving trade or analogous shifts in which new roads can lower deforestation in one location by shifting production to another. This can confound spatially bounded analyses.

One reason for the failure of an empirical literature to converge upon *the* average impact of roads could be that many or most locations do not display the average impact but instead impacts vary. Chomitz’s *Loggerheads* volume, for instance, emphasizes different settings as a core framework.

Bringing this to empirics, with co-authors Weinhold and Pfaff each separately extends in careful empirical work for the Brazilian Amazon an early observation concerning Mexico by Nelson and Hellerstein, who found new roads to increase deforestation less in the places near to prior roads. In the recent Amazon work, prior deforestation is used as the indicator of past development that could affect the impact of new roads. That work is discussed further below but in summary it all suggests that impacts vary importantly with prior development and it supports the Mexico claim.

However, in all of that work data are real constraints, albeit to varying degrees. First, they may not permit the empirics to directly address endogeneity. In this paper, the data permit matching. Second, they tell only part of the story concerning varying impacts, in essence distinguishing only two categories in which impacts differ. Here, we distinguish three and demonstrate why three categories and prior data constraints can explain previous lack of clarity in the literature.

Specifically, this paper brings matching methods to spatially and temporally rich data on roads and deforestation in the Brazilian Amazon. The time period over which the evolution of roads can be tracked and linked to forests is rare, and that much more so with the precision of pixels. This combination allows for the use of lagged roads investments and application of matching. These are significant steps forward in addressing endogeneity in the analysis of infrastructure and these methodological advances are applied to categories expected to show varied impacts.

We find that the impacts of new road investments upon local deforestation are very low where a road was already relatively close by in the past. In addition, we find that the impacts on clearing within the first 5-10 years after a new roads investment is also quite low in pristine areas far from prior roads. It is between these endpoints that new road investments significantly increase rates of deforestation in the next few years. This is where the changes in transport costs are relevant.

The rest of the paper is as follows. Section 2 summarizes prior empirical analysis of road impact as well as the typical von Thunen perspective on road impacts as well as some dynamic models which appear relevant. Section 3 presents the data and the matching methods that we will apply. Section 4 then presents our results and Section 5 summarizes then discusses some implications.

2. Empirical & Theoretical Background

2.1 Relevant Empirics

While many have previously analyzed tropical deforestation, including in the Brazilian Amazon, little empirical work of the sort presented in this paper has been done. A ‘first wave’ of empirics correlated factors of interest with national measures of deforestation.¹ The dominant results are about population density. In cross-country research, it is the most commonly collected variable.

A ‘second wave’ of within-country empirics controls for more variables, with more observations, in testing a given factor.² For the Brazilian Amazon, Almeida 1992 focuses on migration and agricultural colonization. Reis and Margulis 1991 and Reis and Guzman 1992 econometrically find population density, roads and crop area to be important. Pfaff 1999 extended this research, finding population’s spatial distribution (and thus urbanization) to be critical and confirming that roads, not only within the county of interest but also in neighboring counties, are key factors.³

¹ Lugo et al. 1981, Allen & Barnes 1985, Palo et al. 1987, Rudel 1989, Cropper & Griffiths 1994 and Deacon 1994.

² Panoyotou and Sungsuwan 1989 find Thai deforestation driven by population, wood price, income, and Bangkok. Southgate et al. 1991, for Ecuador’s Amazon region, explain population with “prospect of agricultural rents” then explain deforestation with population and other factors. Harrison 1991, for Costa Rica, suggests differing effects of population by region. Kummer 1991 is one of few empirical studies to find only a small role for population growth.

³ At about the same time, and using similar approaches, Chomitz & Gray 1996 analyzed deforestation in Belize and roads, Cropper and Griffiths 1996 examined Thailand, while Nelson and Hellerstein 1997 studied central Mexico and Geoghegan, Wainger and Bockstael 1997 considered the Patuxent Watershed in Baltimore –Washington area.

However, road impacts have not always been so large in such empirics, as noted in Chomitz and Thomas 2003's Amazon analysis. They focused on increasing the variables controlled for and on increasing observations by moving to the level of census tracts from counties. A general note is that biophysical factors such as rainfall were powerful controls, as is more apparent with better data. A road-specific note is that estimated road impacts can be influenced by the level of prior clearing, suggesting attention to prior development in such estimation. As noted above, Nelson and Hellerstein for the Darien region of Panama also find that prior roads influence road impact.

Thus to address these questions it helps to go beyond Chomitz and Thomas 2003's cross-section to track the sequence of road investments over time and relate it to the sequence of deforestation. Pfaff et al. 2007, 2009a and 2009b, using census-tract data like Chomitz and Thomas over time (giving up on the use of census data not available at this level earlier), track this across decades.

These studies examine the tracts receiving road investments (Pfaff et al. 2009a) and, to look for spatial interactions, also at their neighbors not receiving roads investments (Pfaff et al. 2009b). They find that deforestation rises not only in census tracts receiving roads but also in nearby tracts in the same county without roads investments. To the extent permitted by the data, this work takes care in its inference, for instance using lagged roads investments and dummies for counties to control for much unobserved variation over space. This is possible because data for census tracts has over twenty times as many observations as counties (roughly 6000 vs. 300).

That work follows up on Weinhold 2008 (extending Andersen et al. 2002) which also takes care in regressing deforestation on prior changes in roads. These works advanced the idea implicit in Nelson and Hellerstein 1997 by analyzing how road impact differs as a function of context, in particular prior clearing. Estimating an interaction between prior clearing and road investment suggests that prior clearing lowers road impact. Extrapolated, new roads could lower clearing.⁴

Yet the data constraint is critical here, as these works used county data, and in particular the polygons consistent with the county structure over time for which about 250 observations are available. This made it impossible to split the sample or to do anything other than a single interaction term which assumed a monotononic impact of prior development on road impact. While the basic idea of conditional impacts is critical, and supported by their findings of an interaction significantly different from zero, the detailed coefficients require re-examination.⁵

⁴ Why this lowering of deforestation by new roads would occur is not clear in these works. One idea is that there could be spatial intensification in which all development in the region is drawn to the location of the roads, e.g. by migration, and thus clearing in other locations could decrease more than clearing rises where the new roads arise. However, Pfaff et al. 2007's finding with more precise data is that the census tracts in a county receiving a new road increase in deforestation due to the investments in other census tracts. There could be stories concerning trade, such as in New England (see Pfaff 200x and Pfaff and Walker 2009 for discussion of railroad connections to the Midwest and New England reforestation). However, Pfaff and Walker 2009 note this seems not to apply to the Amazon now.

⁵ Another important issue is what the 'average' result is from such a finding. Weighting of observations matters, as using unweighted counties puts a very heavy weight near cities since numerous small counties are found there (this

Pfaff et al. 2009a and 2009b reevaluated these results by examining the much more numerous census tract observations in groups distinguished by level of prior clearing (for instance in 2009a the categories are 0% , 1-50%, 51-75%, >75%). The dominant first two categories show strong increases in deforestation from roads investments. And while the last category is insignificant (with fewer observations), for 50%-75% prior clearing the increase in deforestation resulting from new roads investments is higher, not lower, than it is for the more pristine areas. Thus at the level of census tract data, no deforestation lowering is found and prior development has a non-monotonic influence on the impact of roads, which is lower for very high and low prior clearing. Section 2.2 just below considers reasons why we might expect these empirical results to obtain.

However, despite appropriate care taken in each case above, conditional upon the data's limits, still we are concerned about causal inference given the potential for endogeneity of new roads. Moving to pixel-level data here helps in this by providing sufficient detail to permit matching. Pixel information on distance to the nearest road has other advantages as well, as census tracts are much smaller than counties but not tiny and a new road that crosses a census tract will not have the same impact on all parts of the census tract.⁶ Here, though, we focus on endogeneity.

2.2 Relevant Theory

2.2.1 *Transport Cost (von Thunen⁷)*

A dynamic theoretical model that includes transport cost can address issues of irreversibility and follows the land-use-and-deforestation literature (e.g., Ehui and Hertel 1989, Stavins and Jaffe 1990 and Parks and Hardie 1995). The manager of each hectare j faces a dynamic optimization problem. Risk neutral by assumption, the land manager selects T , the time when land is cleared, in order to maximize the expected present discounted value of returns from the use of hectare j :

$$\text{Max}_T \int_0^T S_{jt} e^{-rt} dt + \int_T^{\infty} R_{jt} e^{-rt} dt - C_T e^{-rt}$$

where:

S_{jt} = expected return to forest uses of the land

appears to be a standard pattern of development; in the Amazon, the rural counties are simply enormous while those near the cities are orders of magnitude smaller). Weinhold gives 'average' claims that would not apply to the region on average in the sense that if you located a road by blindly tossing a dart one would expect those averages to hold. The average of counties (or census tracts for that matter) which applies there would spatially weight by their areas.

⁶ And other supportive evidence concerning varied road impacts is now being generated using pixel data though not using matching to this point (while still using lagged road investments which are a considerable improvement upon contemporaneous investments). Delgado et al. 2009 (following Delgado's master's thesis) supports the idea that the impact of a new road in an already highly developed location may be quite low. Amor and Pfaff 2009 (following Amor's doctoral thesis) supports the idea that a new roads in a pristine area may have a low first-decade impact.

⁷ As in Pfaff 1995, Chomitz and Gray 1996 and Nelson and Hellerstein 1997, this could also be done in a more static fashion but the critical lack remains, ie that after a shock land use adjusts then forest stays constant at that new level.

R_{jt} = expected return to non-forest land uses

C_T = cost of clearing net of obtainable timber value and including lost option value⁸

r = the interest rate

Two conditions are necessary for clearing to occur at time T . First, clearing must be profitable. For clearing to occur, the present discounted rents from non-forest uses will have to more than compensate the manager for the lost returns from forestry uses and the net cost of land clearing:

$$\int_T^{\infty} (R_{jt} - S_{jt}) e^{-rt} dt - C_T > 0$$

However, even if clearing is profitable at time t , it may be more profitable to wait and clear at $t+1$. For example, clearing costs may fall. Thus, the following ‘arbitrage’ condition must hold:

$$R_{jt} - S_{jt} - r_t C_t + \frac{dC_T}{dt} > 0$$

Both conditions must hold for clearing to be preferred. However, if the second-order condition below holds⁹, then either of the necessary conditions is also sufficient for clearing to be chosen.

$$\frac{dR_{jt}}{dt} - \frac{dS_{jt}}{dt} + \frac{d^2}{dt^2} C_t > 0$$

The forest status for each parcel at each point in time will be determined by whether or not these conditions hold or have held previously. Land parcels will have different outcomes across space due to different returns, e.g. from varying land quality and as well as varying access to markets. Returns and costs will also change over time, for both exogenous and endogenous reasons. These all feed through individual decisions to determine aggregate patterns of deforestation over time.

2.2.2 *Partial Adjustment*

The model above assume instantaneous full adjustment on the frontier, i.e. that upon a new road making more locations worth deforesting for agricultural production, for instance, they will all be deforested in the next observed time period. However, adjustment takes time (as noted in various investment literatures) in particular on the frontier. When a road enters a previously less developed or pristine area, the labor and capital required to carry out all of the land-cover change

⁸ A more comprehensive model would also include uncertainty, risk aversion, and forward-looking knowledge of the ability to shift back and forth optimally between cleared and uncleared states. Uncertainty combined with the irreversibilities in deforestation and the ability to learn over time implies an option value to waiting to clear, though as noted neither in the related deforestation literature nor in this paper are individual dynamics the empirical focus.

⁹ For land-use change in a developing country, population and economic growth along with improved infrastructure may lead this to hold, although at a certain stage of development this may be reversed. As development proceeds, environmental protection may become more stringent, returns to ecotourism may well rise, and agriculture can be more capital intensive and require less land. Agricultural returns could fall relative to forest returns on some land.

that may suddenly be economically worthwhile are not present. In particular, a sudden rise in the demand for labor, for instance, will change the marginal cost of hiring additional labor and then what seemed to be worth clearing instantaneously is not because of limited availability of inputs.

Two implications of partial adjustment on the frontier are worth emphasizing. First, given a past investment there will be ongoing clearing of forest in current periods even without more recent transport investments. As inputs become available at the originally envisioned marginal costs of inputs, for instance, the newly profitable locations will be cleared over time. Second, we would expect that the first-decade impact of a new road would be lower where inputs are scarce, e.g. in currently pristine forest areas, relative to where inputs are plentiful. This could even be the case if a quite considerable drop in transport costs has occurred at a location that might remain remote.

2.2.3 Endogenous Development

The model above ignored the possibility that future actions will be influenced by current actions. For instance, upon a road entering an area and connecting it to more developed areas outside of the frontier, other investments that the econometrician does not observe are likely to occur. Settlers will agitate for additional roads investments, for schools, for agricultural subsidies, and more generally for any investment which would improve their quality of life. These are likely to alter the decisions made by other migrants and producers, yielding high spatial path dependence.

The implications of such endogenous development are in a way similar to but are distinct from the possibility of partial adjustment. The similarity is that given past investment such as a new road giving initial access to an area, without other observed new road investments deforestation can be expected to continue as other investments or changes shift the threshold for deforestation.

Endogenous development, though, has an additional implication for the expected impacts from new road investments. It would support the partial adjustment implication that on the far frontier where few observed road investments have occurred road impacts will be low, in this case due to the lack of other, unobserved complementary investments. Further, though, it may suggest that a new road investment in a highly developed location could have low impact. The reason is that if other investments come to dominate the local spatial details of relative net benefits of agriculture then the new road may shift thresholds only inframarginally. Together, the two additional views on the frontier suggest that the transport cost shifts may matter most in an in-between level of prior development, where inputs are available but still other factors do not dominate land use.

3. Data & Methods

3.1 Data

We use highly detailed spatially explicit satellite data on forest overlaid on maps of roads and other factors that influence deforestation. Here, we describe data sources and characteristics.

Regardless of details, however, every variable below is sampled for 100,000 pixels across the Amazon. Considering that the Brazilian Legal Amazon encompasses 5m km², on average our data set has a pixel every 50 km², i.e. a reasonable coverage without points on top of each other.

3.1.1 Deforestation

Deforestation maps were produced in 1997 by IBGE (Instituto Brasileiro de Geografia e Estatística) for their “Diagnostico Ambiental da Amazonia Legal” data product. The pre-1976 clearing is from the RADAM Project vegetation maps, with classes of land cover. The clearing in 1987 is from IBAMA/INPE maps based upon Landsat imagery, as is the 1991 clearing. Our dependent variable for two of our three periods is the difference over time in these measures. For each pixel, we measure deforestation is the pixel was forested in 1976 and cleared in 1987, for our first period, and for our second period if it was forested in 1987 and deforested in 1991.

Our third period is 2000-2004. We have a time gap (i.e., we skip 1991-2000) because we did not want to compare point-in-time forest measures from different sensors to infer pixel-level deforestation. For this period, both years of data are from the PRODES (available from INPE).

3.1.2 Roads

We tracked the evolution of roads over time on maps, providing high spatial specificity. Digital road maps were developed in the Department of Geography at Michigan State University from paper maps by DNER (Departamento Nacional de Estradas de Rodagem), an agency within the Transport Ministry in Brazil. The digital maps that we employ for the analyses in this paper show the distribution of paved and unpaved roads for 1968, 1975, 1985 and 1993.

For these years, for each pixel in our sample and each year in which we have a road map, we measure the distance to the pixel from the nearest paved road and, separately, the distance to the nearest unpaved road. Then we calculated the change, if any, in these nearest-road distances during each period (1968-1975, 1975-1985, 1985-1993). In the measure we are currently using, we assign a zero if the nearest-road distance does not change. We call this “no investment” for that pixel though it really means “no relevant investment”. Alternative measures could reflect the effect of road density, e.g. if another road slightly further away comes in density could matter.

We use these road investments (drops in nearest-road distances) as ‘lagged’ measures. Thus, we will use investment during 1968-1975 to explain 1976-1987 deforestation behaviors and, along the same lines, we will use road investment during 1975-1985 to explain 1987-1991 deforestation and we will use investment during 1985-1993 to explain 2000-2004 deforestation.

It is immediately clear that relationships between road-distance data timing and deforestation timing are not the same in each period. In the third time period, the roads had been in for seven years before our deforestation is measured. In the first and second time periods, the investment was just prior to the measured deforestation but the length of the deforestation measurement period differs and thus the timings are not exactly the same. This is one reason that we start by examining each time period separately. Another is that the true relationship can vary with time.

3.1.3 *Other Drivers*

We integrated maps of ecological conditions and distances to cities to control for their influences in testing our hypotheses about roads. The ecological variables that we can employ here are an index of soil quality, continuous rainfall data¹⁰, and binary variables indicating slope categories (for instance, whether on steeply sloped land or whether on rolling hilly land etc.).

Distance to the nearest city (in km) is computed using both a set of 19 large cities (density over 100 people/km²) and a set of 270 medium and large cities (density over 11 people/km²). These distance variables are used to represent transport costs, to indicate a tract's proximity to a very large city, and to eliminate the census tracts closest to cities from the analysis.

3.2 Matching Methods

3.2.1 *Defining 'Apples to Apples'*

The basic idea of matching to address endogeneity is that the observable differences between locations are central to the non-random allocation of the treatment, in this case the investments in new roads. If they are also central to deforestation, which is the case here, then the observable differences between the treated and untreated locations can bias the coefficient on the treatment. Comparing treated and untreated locations that are as similar as possible in terms of observable factors, i.e. 'apples to apples' comparisons, will then help to reduce the bias in treatment effect.

While those facilitating the application of such techniques agree on this point (Rubin 1980, Rosenbaum and Rubin 1983 and many others), similarity has been defined in different ways. To examine robustness, and to highlight the importance of this definition, we apply two rather different matching estimators. The first is a nearest-neighbors propensity score matching estimator (Rosenbaum and Rubin 1983, Hill et al. 2003), using a fixed number of matched control observations for each treated observation and varying that number from 1 to 50.¹¹ Propensity score matching estimators define similarity based on estimated probabilities of being treated which are generated by a first-stage regression for whether observations are (not) treated.

¹⁰ For discussion of the soil and rainfall data, see Laurance, W.F. *et al.* (2002). *J. Biogeography* **29**, 737.

¹¹ A natural alternative to a fixed number of matched untreated points per treated observation is to fix the window for how good a match needs to be and let the number of matches per point be endogenous to the quality of matches.

The second matching approach we will apply is a nearest-neighbors covariate matching estimator (Abadie and Imbens 2006a) using an inverse weighting matrix to account for the difference in the scale of the covariates. Here again, we employ a fixed number of best matches per treated observation; in Table 2, we apply matching approaches for six matches per treated observation. Covariate matching estimators define similarity without a first-stage regression but rather using the simple distances, in the space of the matching covariates, between the treated and matched.

The computation of standard errors is another difference in how those advancing such techniques have applied them. Abadie and Imbens 2006b show that the common practice of bootstrapping standard errors is invalid with non-smooth, nearest-neighbor estimators such as the propensity score matching estimator with a fixed number of matches that we have chosen to present here (contrasted with kernel versions that assign smoothly declining weights to progressively less-well-matched untreated observations). For propensity score matching, then, we do not bootstrap but follow Hill et al. 2003 in calculating weighted standard errors. For lack of certainty, though, we lean more on the covariate matching standard errors which follow Abadie and Imbens 2006a.

3.2.2 Match Quality

For all of the helpful analytics and important choices described above, the matching approaches control solely for selection on observable factors, unlike using instruments for where policies go. Our experience has been that upon recognizing this, many analysts ask why or if OLS is inferior. The basic reasoning is useful *per se* and it motivates further adjustments in matched estimations.

In short, the attempt to match treated with untreated observations explicitly examines whether in fact there exist untreated observations similar (by whatever criterion) to the treated observations. To the extent that the observed characteristics are not similar in these two groups, OLS uses the information at hand to control for differences but the burden on the specification is considerable:

"Unless the regression equation holds in the region in which observations are lacking, covariance will not remove all the bias, and in practice may remove only a small part of it. Secondly, even if the regression is valid in the no man's land, the standard errors of the adjusted means become large, because the standard error formula in a covariance analysis takes account of the fact that extrapolation is being employed. Consequently the adjusted differences may become insignificant merely because the adjusted comparisons are of low precision. When the groups differ widely in x , these differences imply that the interpretation of an adjusted analysis is speculative rather than soundly based."
(Cochran, in Rubin 1984).

This shortfall of OLS where the untreated are not similar to the treated, however, also relates to matching procedures when match quality is not good. That can motivate using a subset of the

treated observations, and their matches, in particular dropping treated points with poor matches.

Crump et al. 2006 addresses the issue of a lack of covariate overlap, noting that many common estimators become sensitive to the choice of specification (much as Cochran had noted for OLS, and following also related prior work including Heckman, Ichimura and Todd 1997 and 1998). Crump et al. characterize optimal sub-samples for which treatment effects can be estimated most precisely, which under some conditions can be characterized by a rule based on propensity score.

4. Results

4.1 The Importance of The Past

Above the idea of endogenous development was put forward as potentially important for the analysis of impacts of roads investments. Here we start by looking at deforestation rates without current roads investments, recalling that by ‘current’ we mean the immediate lagged time period since we are trying (as have other recent analyses) to avoid problems with road endogeneity.

We consider in Figure 1, for instance, the deforestation rates for locations for which the distance to the nearest paved road did not change during the lagged period. This figure is for investments in roads during 1968 to 1975, the period we use to explain deforestation during 1976 to 1987. Along the diagonal are the locations with the same distance to nearest road in 1968 and 1975. Moving outwards on the diagonal those are locations which were farther from the nearest road.

Thinking in terms of perfect adjustment, we might expect no deforestation at all in these places. The investments were made in 1968 and during the years until 1975 all worthwhile clearing could have already occurred. If we observe all investments, e.g. if roads are the main shift and other factors like slope and rainfall remain constant over time, then forest level may stay fixed.

The perspective of endogenous development, however, would suggest something else entirely. In that view, observed prior road investments would be expected to have led to other investment, potentially unobserved, following upon the prior road investment. In that case, we would expect more of those deforestation-inducing investments where prior roads were relatively close, i.e. to the left within Figure 1. The measured deforestation rates support that this could be going on, as the deforestation rates without roads within 20km of roads is ten times that quite far from roads.

4.2 Non-random Road Investment Can Confound Analysis (even one type of ‘exact’ matching)

The importance of controlling for implications of the past is demonstrated in Figures 2a and 2b. We consider what one might try based upon a belief in perfect adjustment to roads investments. Each ‘row’ contains the estimated impacts of a roads investment, defined as the 1975 distance being at least 5km less than the 1968 distance to nearest road, upon deforestation conditional

upon the 1975 distance to nearest roads being in a given range. That range would be the critical one should it be the case that land use adjusts quickly to road shocks. In fact, if land use were to adjust perfectly we would expect that there would be no effect of investment all across that row because whether an investment occurred during 1968-1975 or not the forest should be the same.

We find several negative and significant impact estimates here, suggesting that having new roads lowers deforestation. This arises even though we are matching for many characteristics of pixels; plus, by working within a row, exclusively we are ‘exact matching’ on the 1975 distance range.

The explanation for these negatives is that we have ignored the past, the 1968 distance which as noted may represent not only past road investment but also followup complementary investment. In our matching, we did not match on the 1968 road distance, i.e. we intentionally ignored the past to implement the model one might typically apply to such data. That drives the negatives.

There are at least two ways to see this. In Figure 2a, the locations with investments are compared to section of the diagonal, which is the locations without investments, solely within that row. The comparison diagonal group is, then, by construction, for a lower 1968 distance than the locations in the rest of the row with the same 1975 distance but higher 1968 distance given the investment. This means that the investments are in places that per endogenous development are less likely to undergo deforestation (as in Figure 1 and 4.1 above), causing roads to appear to lower clearing.

Figure 2b changes the comparison group to the entire diagonal, i.e. full range of 1968 distances. In principle, this could address the spurious negative results (which need not arise and need not be significant; the same figure does not produce as many significant negatives in other periods). However, if the locations with road investments have larger 1968 distances on average than do the places without, again the higher investments are taking place in places with lower clearing. While the spurious negatives are smaller in magnitude, in this case they still signal this issue.

4.3 Regular Matching Plus ‘Exact’ Matching On Prior Development

Given the results in Figures 1 and 2, the results in Figures 3 focus upon controlling for the past. Each column contains the estimated impacts of a roads investment, defined as the 1975 distance being at least 5km less than the 1968 distance to nearest road, on deforestation conditional upon the 1968 distance to nearest roads being in a given range. That range is critical if, as in the view of endogenous development, past early investments lead to relevant unobserved investments.

Here we are doing ‘exact’ matching on the 1968 nearest-road-distance range, which we are doing simply by splitting the sample according to the ranges, as well as standard propensity score (n=4) matching to get apples-to-apples comparisons within each column based on other characteristics. We believe that this should strip out the influence of endogenous development on the baseline,

i.e. the deforestation without investments, so that the variation in impact of road investments can be seen. The way Figure 3a conveys the results shows that prior development affects impact too.

Specifically, we can see that if one starts far enough away from a road in 1968 then even if there is some investment there is not likely to be a significant impact on deforestation (recall that our investment definition allows for relatively small changes in distance to nearest road in 1968-75). That is consistent with both partial adjustment and these investments simply being inframarginal.

Also, we can see the if one starts relatively close to a road in 1968, again having the investment in a new road does not significantly raise deforestation (as we will see in the next section, in all cases these columns are averaging some different realities but still they are worth examination). That is quite consistent with endogenous development predictions that the initial access could set loose a development dynamic distinct from transport costs which could dominate relatively small additional changes in transport cost. Also, the size of the transport cost shift just may be smaller, i.e. the marginal impact of the investment is smaller where there is already road infrastructure.

In the middle columns, though, new road investments clearly raise deforestation significantly, meaning statistically significantly and also in terms of magnitude as a 5% increase across a decade is important. Further, there are no negative results at all even though there are zeros.

This conclusion is echoed in Figures 3b and 3c for the other two time periods, noting that these are both significantly shorter than the initial 1976-1987 time period for deforestation and that on the whole the 1987-1991 clearing rates are lower. In both cases, the only significant results are positive and appear in the middle columns, not the very high nor the very low initial distances. Further, for the third (2000-2004) deforestation period the significant positive results are of the same magnitude as those for the first, longer time period. For the second period the average is clearly smaller though the cells driving that results (as in 4.4 below) contain larger impacts.

4.4 Resurrecting von Thunen | Endogenous Development & Partial Adjustment

Once results in Figures 3 (and 4.3 above) demonstrate controls for the critical prior investments, the next natural question becomes whether with appropriate controls larger investments do more. We know not to look where prior distances to roads are either large or small, for reasons given, however within the 80-300km range of 1968 distance highlighted in Figure 3 we can look harder.

Figure 4 breaks down the columns, in the sense that the cells are compared to column diagonals, i.e. to the locations that do not receive new road investments and have the same 1968 distances. Thus comparisons in columns control for the past and compare larger reductions in transport cost and in each of the three columns in the 80-300km range the impacts in the lower cells are higher. Also it is suggestive that larger investments to the right in the 40-80km row for the 1975 distance yield larger impacts (despite having higher 1968 distances which, as noted, can lower estimates).

5. Discussion

We find that the impacts of new road investments upon local deforestation are very low where a road was already relatively close by in the past. In addition, we find that the impacts on clearing within the first 5-10 years after a new roads investment is also quite low in pristine areas far from prior roads. It is between these endpoints that new road investments significantly increase rates of deforestation in the next few years. This is where the changes in transport costs are relevant.

Upon reflection this is not surprising, given partial adjustment and endogenous development on the frontier (discussed at length in the text). Yet the non-monotonicity has confounded summary. Depending on where new roads are sited, an average could easily find low impact of new roads, even when environmental advocates are right to insist that a given new road will devastate forest. More generally, with varied impacts the data demands are high to identify the impact by category and given the lower-higher-lower impacts as prior development increases any number of claims, including that roads lower deforestation, can erroneously arise using averages or extrapolations.

Non-monotonicity also leads one to consider both temporal dynamics and spatial nonlinearities. From first-decade impacts on deforestation one might claim that a road through pristine areas is less damaging than where prior development was neither high nor low. However, an endogenous development view again is relevant. When follow up investments drive ongoing clearing, e.g. if new roads follow old, long-run pristine impacts are far greater than the immediate forest impact.

Spatially, for instance thinking of species habitat as a ‘co-benefit’ of actions paid for by carbon, as demonstrated in Amor 2008 the habitat fragmentation by roads in pristine areas is enormous. Thus even in the short run, e.g. the first decade, roads through pristine areas look less positive.

The policy implication of the correct results is important. It seems clear from these results that a nation could choose to achieve its development goals with different shapes of transport networks and that which shape is chosen will affect deforestation. These choices could earn global carbon payments. Further, the same logic could in principle apply to pipelines too (Feder(?) et al. 2008).


The idea that variations in impact or deviations from the average impact are important for policy choice is generally relevant for forest. If protected area impacts on deforestation vary depending upon the context, then policy makers locating a new protected area can target based upon impact (see, e.g., Andam et al. 2008 and Pfaff et al. 2009). The same holds true for any type of payments for ecosystem services (see, e.g., Pfaff et al. 2008 and Robalino et al. 2008). Targeting matters.

For both development and conservation policies, further research in other locations would help to establish how general are these results. While the Brazilian Amazon is one critical specific area, and also one for which excellent data exist, to the extent possible replication would be of value.

Figure 1

**measured 1976-1987 deforestation rates (fraction cleared)
as a function of the 1968 and 1975 distances to nearest road**

Distance to road 1975 (km)




Distance to road 1968 (km)

Figure 2a
measured road-investment impacts, using matching,
explaining 1976-1987 deforestation rates (fraction cleared)
“by row”, i.e. by end-of-period (1975) distance to the nearest road

(“road investment”, i.e. 1975 distance is at least 5km lower than 1968 distance, is compared with “no road investment”, i.e. the diagonal within the same row, with Matching (propensity score, n=4) for apples-to-apples comparison in row)

Distance to road 1975 (km)



Distance to road 1968 (km)

Figure 2b
measured road-investment impacts, using matching,
explaining 1976-1987 deforestation rates (fraction cleared)
“by row”, i.e. by end-of-period (1975) distance to the nearest road

(“road investment”, i.e. 1975 distance is at least 5km lower than 1968 distance, is compared with “no road investment”, i.e. the diagonal, but here the entire diagonal, with Matching (propensity score, n=4) for apples-to-apples comparison to the untreated)




Figure 3a

**measured road-investment impacts, using matching,
explaining 1976-1987 deforestation rates (fraction cleared)**

“by column”, i.e. by start-of-period (1968) distance to the nearest road

*(“road investment”, i.e. 1975 distance is at least 5km lower than 1968 distance,
is compared with “no road investment”, i.e. the diagonal within the same column,
with Matching (propensity score, n=4) for apples-to-apples comparison in column)*

Distance to road 1975 (km)





Figure 3b
measured road-investment impacts, using matching,
explaining 1987-1991 deforestation rates (fraction cleared)
“by column”, i.e. by start-of-period (1975) distance to the nearest road

(“road investment”, i.e. 1985 distance is at least 5km lower than 1975 distance, is compared with “no road investment”, i.e. the diagonal within the same column, with Matching (propensity score, n=4) for apples-to-apples comparison in column)

Distance to road 1985 (km)



Distance to road 1975 (km)

Figure 3c

**measured road-investment impacts, using matching,
explaining 2000-2004 deforestation rates (fraction cleared)
“by column”, i.e. by start-of-period (1985) distance to the nearest road**

*(“road investment”, i.e. 1993 distance is at least 5km lower than 1985 distance,
is compared with “no road investment”, i.e. the diagonal within the same column,
with Matching (propensity score, n=4) for apples-to-apples comparison in column)*




Figure 4
measured road-investment impacts, using matching,
explaining 1976-1987 deforestation rates (fraction cleared)
“by cell”, i.e. by start-and-end-of-period (1968 & 1975) distances

(“road investment”, i.e. 1975 distance is at least 5km lower than 1968 distance, is compared with “no road investment”, i.e. the diagonal within the same column, with Matching (propensity score, n=4) for apples-to-apples comparisons of cells)

Distance to road 1975 (km)

